

EXPRESS MAIL NO. EK927915228US

Docket No. 00AB025

Ref. No. 110003.97001

PATENT APPLICATION FOR
GLOBAL RESOURCE LOCATOR FOR
AUTONOMOUS COOPERATIVE CONTROL SYSTEMS

By

Francisco P. Maturana
Sivaram Balasubramanian

Express Mail Mailing Label Number #EK927915228US

Date of Deposit July 24, 2000

Pursuant to 37 C.F.R. § 1.10, I certify that I am personally depositing this paper or fee with the "Express Mail Post Office to Addressee" service of the United States Postal Service on the above date in a sealed envelope (a) having the above-numbered Express Mail label and sufficient postage affixed, and (b) addressed to: BOX PATENT APPLICATION, Commissioner of Patents and Trademarks, Washington, DC 20231.

Date: July 24, 2000


Ellen T. Roy

GLOBAL RESOURCE LOCATOR FOR
AUTONOMOUS COOPERATIVE CONTROL SYSTEMS

CROSS-REFERENCE TO RELATED APPLICATIONS

--

5

STATEMENT REGARDING FEDERALLY
SPONSORED RESEARCH OR DEVELOPMENT

--

BACKGROUND OF THE INVENTION

10 The present invention relates to industrial control systems and the like and in particular to an industrial control system comprised of many distributed controllers dividing job tasks among themselves using "market model" based bids and counter bids.

15 In a centralized model for industrial control, a single central controller (being a specialized computer) coordinates all aspects of the controlled process. Input signals representing values from sensors on the controlled process are communicated to the central controller which executes a control program to develop output signals which are communicated to actuators on the controlled process.

20 The centralized model is conceptually simple and requires programming only a single device. Nevertheless, the centralized model has some disadvantages. Control systems using this model are particularly vulnerable to failure of the single central controller. Further, communication of all I/O signals to a central location and executing the control program on a single computer, particularly for large control systems, can place undue burdens on processing resources and communication bandwidth.

25 In the distributed model for industrial control, the control program is executed by a number of spatially separate controllers intercommunicating only as needed on a common network. By placing these controllers near relevant I/O points, the communication of large amounts of I/O data is diminished. Having multiple control devices can also reduce the susceptibility of the control system to failure of any one device.

30 One difficulty with distributed control is that of developing the multiple control programs to dividing the control tasks among the various distributed controllers and coordinating their actions. One promising method of both dividing the control task

among the controllers and coordinating their actions borrows from a market model of the economy in which many different autonomous individuals organize themselves (through bidding and counter bidding) to produce complex products or services without central control. In such autonomous cooperative systems (ACS), a job description is presented to a large number of autonomous cooperative units (ACUs), which, based on knowledge of their own capabilities and limitations, bid on portions of the job and counter bid in response to requests for bids. Control programs are simply developed and the potential exists for control programs to be quickly changed as new circumstances develop. Examples of this would be if it is desired to produce a new product or if ACUs (and associated equipment) are introduced or removed from the system. A description of some such of autonomous control systems are described in co-pending patents.

For an autonomous cooperative system to produce an optimized outcome, it is desirable that many different divisions of the job among the ACUs be explored and thus that a large number of bids be collected from many different ACUs. This is done by allowing each ACU to make multiple simultaneous requests for bids from other ACUs. It follows that a single ACUs may receive multiple requests for bids from multiple other ACUs each of which must be considered in a separate "context" and each of which may mature into a different completed bid.

The bidding process for each context requires a large number of bidding messages be transmitted over the network. Such messages generally include "bid requests" and replies to bid requests including "bid success", "bid failure" and "counter bid" messages. The numbers of messages increases geometrically with the number of ACUs involved and these messages can easily overtax even high-speed communication networks thus delaying the determination of a control solution. This delay limits the usefulness of an ACS in responding quickly to changing control situations and effectively limits the size of control problems that may be advantageously handled by the ACS.

What is needed is a way to realize a freely scalable ACS system that may rapidly determine an optimized control solution that may be implemented with existing hardware having network bandwidth and processing limitations.

BRIEF SUMMARY OF THE INVENTION

The present invention provides an ACS significantly reducing the number of bid messages needed to provide a solution, and thus requiring less network bandwidth and ACU processing power. The reduced number of bid messages is the result of two

techniques. The first technique is the creation of a registry that directs a given ACU wanting to request a bid to other likely ACUs for receiving that bid. Bid messages sent to ACUs having a low likelihood of successfully responding are thus reduced.

In the preferred embodiment, at the time the ACUs are initialized, each ACU transmits its capabilities (such as might be provided to a requesting ACU) one time only to the central registry rather than many times to each requesting ACU. The central registry, by providing a single source overview of the capabilities of other ACUs, can significantly shorten the time required for a given ACU to find other ACUs of particular capabilities. The registry thus acts like a central clearinghouse or central market in simplifying transactions.

The second technique of the present invention employs a cache-like structure in each ACU into which portions of the registry may be downloaded. Unlike a conventional computer cache, the cache holds potential relationships between the given ACU and its bidding partners. This technique reduces network traffic by allowing the ACUs, in many cases, to accurately identify other relevant ACUs without referring directly to the registry eliminating the registry directed messages. When the necessary data is not found at the cache, a request may be made of the registry and the registry may use this to further refine its prediction of which data should go into the cache.

Specifically then, the present invention provides a method of coordinating a plurality of autonomous cooperative units implemented in computers intercommunicating on electronic network and operating to divide a predefined job among ACUs by a bidding process including bids communicating among the ACUs. The method provides that for each ACU, a subset of all ACUs of the system are identified as potential bid request receivers. This subset may initially be based on the capabilities of the ACUs as being compatible with capabilities of other ACUs. This subset is stored as a list of potential bid receivers in the computer implementing the ACU. When the ACU receives a description of the job, it preferentially communicates requests for bids related to the description of the job from the given ACU only to other ACUs on the list of potential bid request receivers.

It is thus one object of the invention to provide a systematic exploration of solution space as regards possible combinations of ACUs to compete bids for a given job. The list is selected to indicate likely candidates for the bidding process thus greatly simplifying the number of ACUs that need to communicate and reducing total communications.

The method may include storing at a designated computer a relation table relating each ACU to capabilities of potential bid request receivers for that ACU. The designated computer may collect from each ACU on the network, that ACU's capabilities and may reply to requests from these ACUs, looking for potential bid request receivers, with those
5 ACUs having capabilities matching the capabilities of potential bid request receivers of the given ACU in the relation table.

Thus it is another object of the invention to centralize communication of capabilities among ACUs and thus to eliminate many separate transactions between ACUs that would redundantly communicate these capabilities between ACUs.

10 Each ACU may transmit its capabilities to the designated computer upon initialization of the ACU.

Thus it is another object of the invention to automatically generate the relation table at the initialization thus eliminating it as a task necessary to the programmer and reducing the burden of this process during run time.

15 The designated computer may also implement at least one of the ACUs.

Thus it is another object of the invention to provide the benefit of the designated computer without departing from the distributed model of the autonomous cooperative system.

The process of preferentially communicating requests for bids with other ACUs
20 based on this list may include the steps of reviewing the list of potential bid request receivers for ACUs likely to complete a portion of the job based on information about likelihoods incorporated into the list. When there are no likely ACUs, the method may include communicating with the designated computer to receive a new cluster of ACUs to update the list of potential bid request receivers and then preferentially communicating
25 with an ACU on the updated list.

Thus it is another object of the invention to provide an efficient communication between the more comprehensive list of potential bid request receivers contained at the designated computer and the smaller but more efficient list contained at the device implementing the ACU. This updating process potentially allows each ACU to have
30 ultimately the full range of ACUs in front of it but allows internal and efficient access to the most likely ACUs within that range.

When the ACUs in the list are considered unlikely to successfully complete a portion of the project because of their historical responses, the designated computer provides a new list of ACUs that are simply different from those previously provided.

Where the likelihood is based on the lack of ACUs of a given capability (for example, as a result of an unexpected job) the designated computer alters relation list to provide ACUs under the criterion of having the needed capability.

Thus it is another object of the invention to exploit whatever information is
5 available to focus the bidding process on the most likely bidding candidates.

The foregoing and other objects and advantages of the invention will appear from the following description. In this description, reference is made to the accompanying drawings, which form a part hereof, and in which there is shown by way of illustration, a preferred embodiment of the invention. Such embodiment and its particular objects and
10 advantages do not define the scope of the invention, however, and reference must be made therefore to the claims for interpreting the scope of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a simplified perspective view of a manufacturing process including a number of separate machines intercommunicating via distributed computers connected
15 over a network such as may be used with the present invention;

Fig. 2 is a block diagram of one of the computers of Fig. 1 showing a standard architecture for control including a processor, a memory holding data and programs executed by the processor, a network card providing an interface to the network of Fig. 1 and I/O circuits for communicating with the machines of Fig. 1;

Fig. 3 is a functional diagram of the program and data structures stored in the memory of Fig. 2 including programs implementing autonomous cooperative units, subcontract bid managers associated with each autonomous cooperative unit, a cache-type local resource locator and optionally a global resource locator, the latter which may be implemented in one but not all of the computers;

Fig. 4 is a hierarchical diagram showing organization and communications between the programs and data structures of Fig. 3 as implemented in different computers of Fig. 1;

Fig. 5 is a block diagram showing the operation of an autonomous cooperative unit its subcontract bid manager, the local resource locator and global resource locator of
30 Fig. 4 when a new autonomous control unit is initialized such as creates lists of potential bid request receivers in the local resource locator;

Fig. 6 is an expanded fragmentary representation of data stored in the global resource locator in a registry and relation table, a portion of which as is communicated

with and stored in the local resource locator the latter which also holds a success history table;

Fig. 7 is a flow chart showing programs of the subcontract bid managers in managing the data of the global resource locator and the local resource locator;

Fig. 8 is a flow chart of the operation of the subcontract manager of Fig. 4 in updating the success history of Fig. 6;

Fig. 9 is a tree structure showing a virtual cluster assembled by operation of the global resource locator and local resource locator providing for success weights from the success history table indicating methods of calculating likely success in bid

subcontracting;

Fig. 10 is a diagram representing the timing constraints on bidding in the present invention; and

Fig. 11 is a representation of the data structures and functions collected in the global resource locator of Fig. 3.

DETAILED DESCRIPTION OF THE INVENTION

Referring now to Fig. 1, a distributed industrial control system 10 may include a plurality of machines 12, for example, manufacturing machines such as drills, lathes, ovens, mills and the like, each associated with an electronic computer 14 configured for electrical communication through I/O lines 30 with the machines 12.

The electronic computers 14 may be linked to each other via a network 16 of a type well known in the art allowing for connected messaging or other communication protocol between the various computers 14 whereby each of the computers 14 may produce messages for or consume messages from other computers 14. A human machine interface HMI 20, being a conventional computer terminal or other similar device, may be attached to the network or one of the computers (as shown) allow for programming of the various computers 14 or data entry as will be described.

Referring now to Fig. 2, each computer 14 includes a network interface 18 of conventional design for transmitting and receiving messages on the network 16 and communicating them to an internal bus 22. The internal bus 22 links the network interface 18 with computer memory 24, a processor 26 and I/O circuits 28, the latter which provide I/O lines 30 leading to sensors or actuators on the machines 12. A secondary communication port 32 may be available for connection to the human machine interface 20 as described above. The memory 24 may hold within it a conventional

multitasking operating system, for example, Windows NT (not shown) under which the various programs of the present invention may be simultaneously executed as tasks.

Referring now to Fig. 3, as is generally understood in the art and as is described in the above referenced patents, under the operating system, in each computer 14, the present invention executes tasks that implement autonomous cooperative units (ACUs) 34 configured for particular machines 12 with which the computer 14 and ACU 34 are associated. As part of its creation, each ACU 34 is programmed, for example, through HMI 20, with data representing the capabilities 48 of the machine 12 with which it is associated. This data indicates both the generic operation or operations performed by the machine 12 (termed: "services") and particular values quantitatively delineating that service (termed: "parameters"). For example, a rolling mill would have a service of "rolling material to reduced thickness" and might have parameters indicating maximum rolling speed, percentage reduction and the like. The particular capabilities 48 will be defined for specific machines 12 according to a standard job description language (JDL).

Each ACU 34 also includes a set of goals indicating a local optimization that the ACU 34 will try to achieve within its parameters, and programs executing protocols for receiving the description of the job (written in JDL) that the ACUs 34 are called upon to complete.

Upon receipt of the job description, each ACU 34 will parse the description, bidding on portions of it matching their capabilities 48 and forwarding other portions in a subcontracting process to other ACUs 34 along with requests for bid messages. Preferably, the ACUs 34 communicate with other ACUs 34 via an agent language such as KQML. ACUs 34 may successfully bid on parts of the job reply with a bid success message upward to the bid requesters and in this way complete bids on the total job may be assembled at the uppermost ACU 34 in the bidding chain. Complete bids, meaning those having one ACU 34 successfully bidding on each part of the job, are compared to select a winning bid describing how the job will be allocated among the ACUs 34.

Through this bidding process, the ACUs 34 mimic the actions of individuals within a market economy to divide up the parts of a job for execution by the various ACUs 34 without need for centralized coordination. When the machines 12 are reconfigured (e.g., configured to different settings, added or removed), the new set of ACUs 34 may quickly reallocate their responsibilities using the same processes.

Referring still to Fig. 3, the present invention adds to the prior art ACUs 34 contained within the computers 14 three additional structures.

The first additional structure is a subcontract bid manager (SCBM) 36. The SCBM 36, which will be described in greater detail below, generally manages requests for bids produced by the ACUs 34. In this way, the operation of the SCBM 36 in steering bids (as will be described) is invisible to the ACU's, simplifying their design. A given computer 14 may contain a number of ACUs 34 each associated with a SCBM 36. The SCBMs 36 may communicate directly with other ACUs 34 in the particular computer 14 or another computer 14 and may also communicate with the local resource locator (LRL) 38 also contained in the computers 14 which serves to direct them to likely ACUs 34 to receive requests for a bid.

The second additional structure is this LRL 38 which is replicated only once for each computer 14 and which provides a cache-like storage area shared by the SCBMs 36 and communicating with the SCBMs 36 on the high speed internal bus 22 as opposed to the slower network 16. Operation of the LRL 38 will also be described in greater detail below.

The third additional structure is a global resource locator (GRL) 40. The GRL 40 is generally implemented in a single of the computers 14 as indicated by the dotted lines however its influence extends among all of the computers 14. As will be described in more detail below, and referring also to Fig. 11, the GRL 40 provides a number of functions including serving as a central clearinghouse for information needed by the ACUs 34 in the bidding process per registry table 51. This information may be selectively downloaded to the LRL 38, based on the operation of a capability matcher 108, to create communication clusters which serves to reduce communications over network 16. The GRL 40 also modifies these clusters according to learning protocols 104, and applies coordination rules (such as implements bid expiration timers) per context coordination protocols 102. The GRL 40 is also an ACU 34 and generally handles inter ACU communication per communication protocols 106.

The following processes are implemented by data and programs implemented in these additional structures with little additional programming of the ACU's 34. Although particular functions will be ascribed to particular ones of these structures of the GRL 40 the LRL 38 and the SCBM 36, it will be understood to one of ordinary skill in the art that the function responsibilities can in some cases be moved from one element to another and that the invention should not be considered to be limited to this functional division except insofar as it is essential to the purposes described herein.

Referring now to Fig. 4, in an example control system 10 having multiple computers 14a-14c, the GRL 40 may be implemented in the first computer 14a to serve in shared capacity with each of computers 14a through 14c. Conversely, each computer 14a to 14c has a separate LRL 38a through 38c and separate ACUs 34a and 34a' (in computer 14a) ACU 34b (in computer 14b) and ACU 34c (in computer 14b). Each ACU 34a, 34a', 34b and 34c are associated with corresponding SCBMs 36a-36c.

Referring now to Figs. 1, 4 and 5, in initializing the control system 10, the ACUs 34a through 34c are created, for example, through commands and data entered into the HMI 20, according to a prewritten ACU template and the entering of data configuring each ACU's particular services, parameters and goals as indicated by process block 41.

As indicated by subsequent process block 42, at the time of this initialization, each ACU 34 sends a registration message to the LRL 38 of its computer 14. The message describes the capabilities of the ACU 34 including both its services and parameters. The LRL 38 uses this information to create a row entry 44 in local directory 43 as indicated in Fig. 6. The row entry 44 includes an identification 46 of the ACU 34 in a first column and capabilities 48 of the ACU 34 in a third and fourth column. The identification 46 may simply be the physical address of the ACU 34 on the network allowing ready communication with the particular ACU 34. The second column is reserved for listing related ACUs 34 as will be described and these ACUs 34 may also be identified by their physical address on the network 16.

The LRL 38 acknowledges this message and the ACU 34 waits for an incoming job description language message to begin the bidding process as described above.

Next, as indicated by process block 50, the LRL 38 registers the information provided by the ACU 34 with the GRL 40. Because the LRL 38a in this example is shared with two ACUs 34a and 34a' in computer 14a for this LRL 38a, there will be a second row entry 44' for the second ACU 34. In this example, the ACU 34a is associated with an oven and has as its service, heating and parameters X_1 , Y_1 , X_3 .

In the registration process of process block 50, the LRL 38 sends to the GRL 40 the identification 46 and the capabilities 48 of each row entry 44 and 44'. As noted before, service information indicates generally the type of operation in the manufacturing environment performed by the machine 12 associated with the ACU 34.

The GRL 40 places this material in a registry table 51 as row entries 45. At this time the GRL 40 may create a relation table 54 indicating each registered ACU 34 by its identification 46 and in a second column indicating dependent services 56 likely to be

51, the GRL 40 scans the relation table 54 and determines that the ACU number 1 may logically send requests to bid to this ACU number 3. The data for ACU number 3 (a potential bid request receiver) is then enrolled in the LRL 38 responsible for ACU number 1 beneath the row entry 44 for ACU number 1. In this way, an individual ACU 34 may consult with the LRL 38 to find likely candidates for receiving bid requests without network traffic.

At the same time that the potential bid request receivers are enrolled in the LRL 38, they are enrolled in success history table 72 also held in the LRL 38. A first column of the success history table 72 holds the identification 46 of an ACU 34 associated with the LRL 38 and a second column holds the identification 46' of the potential bid request receiver downloaded to the LRL 38. A third column shows a priority values 78 indicating likely success in bids between these ACU's based on historical data to be collected. Initially this priority value is set to one however it will vary over time as more historical data is collected as to how likely it is that the indicated ACU will return a successful bid.

At the conclusion of this registration process, each of the ACUs 34a through 34c is registered with the GRL 40 and the associated LRLs 38a through 38c contain a list of potential bid recipients for their ACUs 34.

Referring now to Fig. 7, the value of the LRL 38 and the GRL 40 is illustrated as a given ACU 34 receives a bid request associated with a job description or portion of the job description (in JDL) as indicated by process block 64. According to standard operation of the ACU 34, the ACU 34 examines the job description to see if it has parts that match the capabilities of the equipment 12 associated with the ACU 34 as indicated by decision block 66. If the ACU cannot perform a part of the job, the bid is rejected and communicated to the requestor as indicated by arrow 67. On the other hand, if the ACU 34 can perform a part of the job, the ACU 34 consults with its LRL 38 via the SCBM 36 as indicated by decision block 68 to see whether they are potential candidates for sub bidding.

At process block 64, it should be noted that multiple ACUs may send bid requests to a particular ACU 34 and each is provided with a separate context. Multiple contexts allow a given ACU to act as if the requests for bids were received by separate ACUs with the same address. Referring again to Fig. 6, this process involves reviewing the rows of the LRL 38 beneath the inquiring ACU's row entry 44 to see if there are other ACUs 34 enrolled with the necessary capabilities required by parts of the job description not being performed by the inquiring ACU 34. If no potential bid requesters having suitable

capabilities are listed in the LRL 38, the process continues at process block 70 as will be described below.

On the other hand, if potential bid requesters are listed in the LRL 38, then the success history table 72 is consulted to determine whether these candidates have
5 historically provided successful bid responses. The identification 46 of the particular ACU 34 seeking to make a bid request is located in the first column of the success history table 72 and the identification 46 of the potential bid request receiver is located in a second column together to define a row which includes a priority value 78. This priority is compared against a predetermined threshold, for example, 0.5 and if it is greater than
10 that, a bid request receiver is considered successfully found and a request for bid will be forwarded to this identified ACU per process block 94 as will be described below. If however, none of the candidates subcontracting ACUs have sufficiently high priority, the program proceeds again to process block 70. Repeated failure to find successful bid recipients in the success history table 72 may be used to modify the threshold against
15 which priorities are compared.

Failure to find a suitable bid recipient in the LRL 38 is not conclusive, because LRLs 38 tend to be limited by the small size of the computers 14 and far greater information is contained in the registry table 51 of the GRL 40. Accordingly, if no suitable bid request receiving ACU is found, the SCBM 36 communicates with the GRL
20 40 to obtain a new cluster of related ACUs 34. The GRL 40 may respond in two ways according to the learning protocol 104 (shown in Fig. 11). If the needed bid recipient ACU is of dependent services 56 not listed in the relation table 54 for the requesting ACU 34, the GRL 40 will update the relation table 54 by adding the new dependent service 56 and using this new dependent service 56 will rescan through the registry table 51 to
25 obtain a new cluster of ACUs.

Alternatively, it may be the case that process block 70 is reached when bid request receiving ACUs with the proper dependent capability were found in the LRL 38 but their priority values 78 in the success history table are too low. In this case, the GRL 40 will select a new set of ACUs from the registry table 51 different from those listed in the
30 record 86.

Tracking of previously provided ACUs may be held in the record 86 so that rejected ACUs will not be resubmitted to the LRL 38 until all others have been exhausted.

ACUs B and C for their relationships with ACU's D and E and F and G, respectively. By following the path of highest priorities (implicit in the process described above), quicker conclusions of a bid will be obtained (with fewer bid messages) so long as the bid successes are consistent on a historical basis.

Because of the extremely large solution space and large number of possible chains through ACUs in the bidding process, it is desirable to limit the amount of time before which a response to a request for bid must be received. This prevents lost messages from stalling the process and truncates extremely long searches through possible solution space.

Referring now to Fig. 10, accordingly each request for bid is attached to a response time value 92. This is done by the global resource locator 40 as it handles communications of messages related to different contexts per the context coordination protocols 102 shown in Fig. 11. As noted above, a request for bid will include the portion of the description that could not be implemented by the current ACU and thus requires further requests for bid from other ACUs. This response time value 92 is propagated in modified form in all subsequent request for bids by those ACUs 34 receiving the initial request for bid. The modification subtracts from the response time value 92 at each level in the bidding process, time necessary for processing the initial request for bid. So for example, and ACU A may transmit a request for bid 91 having a response time value 92 of 1.0 to ACU B. ACU B in turn may make other requests for bids 91' having attached response time values 92 of 0.7, allowing 0.3 of processing time for ACU B to respond. This chain is continued with each subsequent ACU 34 requesting a response time value that is less than the time allotted to it from the previous ACU 34 so that prompt response may be guaranteed. The response time values 92 of Fig. 10 are preserved independently for each context of the ACU.

In this way, it can be assured that ACU B receives its bid responses in sufficient time to forward them to ACU A in the time it requires. Bids received after the response time value 92 are treated as failures.

It is specifically intended that the present invention not be limited to the embodiments and illustrations contained herein, but modified forms of those embodiments including portions of the embodiments and combinations of elements of different embodiments as come within the scope of the following claims.